PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04Q 11/04, H04L 12/56

A1

(11) International Publication Number: WO 00/36871

(43) International Publication Date: 22 June 2000 (22.06.00)

(21) International Application Number: PCT/SE99/02278

(22) International Filing Date: 7 December 1999 (07.12.99)

(30) Priority Data:
9804320–1 15 December 1998 (15.12.98) SE

(71) Applicant *(for all designated States except US)*: TELIA AB (publ) [SE/SE]; Mårbackagatan 11, S–123 86 Farsta (SE).

(72) Inventor; and
(75) Inventor/Applicant *(for US only)*: KAVAK, Nail [SE/SE]; Myrstuguvägen 359, S–143 32 Vårby (SE).

(74) Agent: PRAGSTEN, Rolf; Telia Research AB, Vitsandsgatan 9, S–123 86 Farsta (SE).

(81) Designated States: EE, JP, LT, LV, NO, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published
With international search report.
Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: DATA TRANSMISSION SYSTEM ADAPTED TO PROVIDE INTERWORKING BETWEEN RSVP AND MPLS

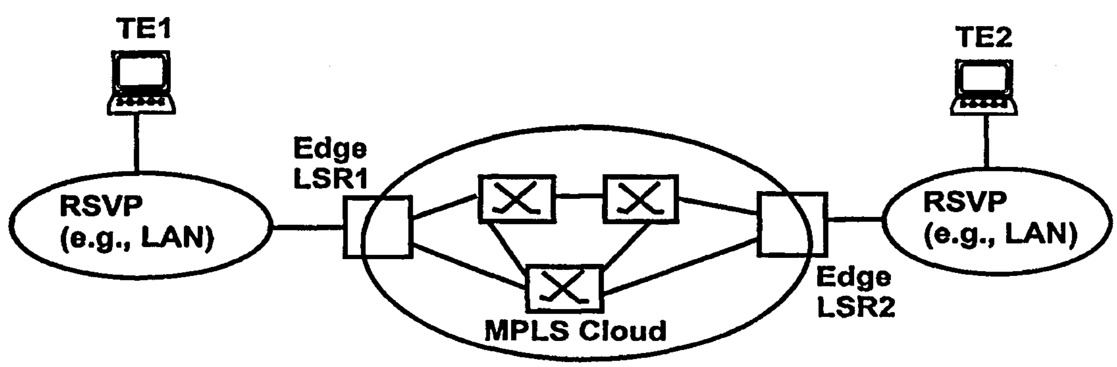

(57) Abstract

The invention provides a data transmission system which includes a backbone network, a plurality of routers, and a plurality of end user equipments and which provides carrier scale end–to–end Quality of Service (QoS) through interworking between RSVP and MPLS. In particular, the system establishes end–to–end data transmission paths over the backbone network, end user equipments use RSVP messages to convey their specific requirements and the system uses MPLS to intercept the RSVP messages and establish Label Switched Paths (LSPs) over the backbone network, based on the content of the RSVP messages.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	**ES**	Spain	**LS**	Lesotho	**SI**	Slovenia
AM	Armenia	**FI**	Finland	**LT**	Lithuania	**SK**	Slovakia
AT	Austria	**FR**	France	**LU**	Luxembourg	**SN**	Senegal
AU	Australia	**GA**	Gabon	**LV**	Latvia	**SZ**	Swaziland
AZ	Azerbaijan	**GB**	United Kingdom	**MC**	Monaco	**TD**	Chad
BA	Bosnia and Herzegovina	**GE**	Georgia	**MD**	Republic of Moldova	**TG**	Togo
BB	Barbados	**GH**	Ghana	**MG**	Madagascar	**TJ**	Tajikistan
BE	Belgium	**GN**	Guinea	**MK**	The former Yugoslav Republic of Macedonia	**TM**	Turkmenistan
BF	Burkina Faso	**GR**	Greece			**TR**	Turkey
BG	Bulgaria	**HU**	Hungary	**ML**	Mali	**TT**	Trinidad and Tobago
BJ	Benin	**IE**	Ireland	**MN**	Mongolia	**UA**	Ukraine
BR	Brazil	**IL**	Israel	**MR**	Mauritania	**UG**	Uganda
BY	Belarus	**IS**	Iceland	**MW**	Malawi	**US**	United States of America
CA	Canada	**IT**	Italy	**MX**	Mexico	**UZ**	Uzbekistan
CF	Central African Republic	**JP**	Japan	**NE**	Niger	**VN**	Viet Nam
CG	Congo	**KE**	Kenya	**NL**	Netherlands	**YU**	Yugoslavia
CH	Switzerland	**KG**	Kyrgyzstan	**NO**	Norway	**ZW**	Zimbabwe
CI	Côte d'Ivoire	**KP**	Democratic People's Republic of Korea	**NZ**	New Zealand		
CM	Cameroon			**PL**	Poland		
CN	China	**KR**	Republic of Korea	**PT**	Portugal		
CU	Cuba	**KZ**	Kazakstan	**RO**	Romania		
CZ	Czech Republic	**LC**	Saint Lucia	**RU**	Russian Federation		
DE	Germany	**LI**	Liechtenstein	**SD**	Sudan		
DK	Denmark	**LK**	Sri Lanka	**SE**	Sweden		
EE	Estonia	**LR**	Liberia	**SG**	Singapore		

DATA TRANSMISSION SYSTEM ADAPTED TO PROVIDE INTERWORKING BETWEEN RSVP AND MPLS

The present invention relates to a data transmission system adapted to provide interworking between a Resource Reservation Protocol (RSVP) and Multi Protocol Label Swapping (MPLS) to provide a carrier scale end-to-end Quality of Service (QoS) architecture, and a method for establishing end-to-end data transmission paths in a data transmission system according to the invention.

RSVP has been specified by the Internet Engineering Task Force (IETF). It specifies a mechanism for the allocation of resources in connectionless Internets. If RSVP is to work optimally, the underlying network must also provide, besides 'best effort', different service classes (for example, ATM). RSVP may be used over ATM as a means of offering secure service to application users. The RSVP, in this case, uses resources on a packet/application level which are later converted, or mapped, to guarantees on the ATM level. There exists within ATM, certain solutions which describe how the mapping of RSVP over ATM can be achieved. All of these solutions use the classic scheme for IP over ATM, (i.e. Hop-by-hop).

It is widely believed that the RSVP has scalability problems in wide area networks due to the memory consumption and processing power requirements. At the present time, RSVP is the only mechanism through which applications may transmit their QoS (Quality of Service) requirements.

MPLS (Multi Protocol Label Swapping) is intended to be used in backbone networks and scales better than RSVP. However, MPLS is not available to a terminal connections that is not end-to-end. Also, mechanisms are not available to convey a user's demands to MPLS.

In a data transmission system, in accordance with the present invention, end-users convey their QoS requirements via RSVP to MPLS which aggregates the relatively small individual RSVP flows into larger flows in one, or a number of, data flows. The advantages to be gained from this are that there is no need for a public network to remember small flows and that such networks have more details, or more exact information, concerning a user's QoS requirements. The effect is that

demands on transmission capacity and network resources, for example, processor and memory, are reduced. RSVP control packages are tunnelled via an MPLS network to make the network transparent to RSVP. Terminals use RSVP to transmit their demands on resources to the network. Access nodes have MPLS functionality which allows them to aggregate data flows. In addition, many small data flows are transported over only one connection if they have the same QoS characteristics.

According to a first aspect of the present invention, there is provided, a data transmission system including a backbone network, a plurality of routers, and a plurality of end user equipments, said system being adapted to establish end-to-end data transmission paths over the backbone network, characterised in that said end user equipments are adapted to use RSVP messages to convey their specific requirements and in that said system is adapted to use MPLS to intercept said RSVP messages and establish Label Switched Paths (LSPs) over said backbone network, based on the content of said RSVP messages.

The system may be adapted to provide end-to-end Quality of Service (QoS) through interworking between RSVP and MPLS.

The system may be adapted to aggregate a number of RSVP messages, having the same QoS semantics, into a single LSP.

The system may further include a plurality of end user networks, to which said end user equipments are adapted to be connected, and a plurality of routers adapted to connect said end user networks to said backbone network.

At least some of said end user networks may be Local Area Networks (LANs).

The backbone network may be an ATM network, adapted for the transmission of IP data packets and including a number of interconnected ATM switches, at least two of said routers may be edge routers for said ATM network, and the edge routers may be Label Switched Routers (LSRs), adapted to be connected to end user equipments for the establishment of unidirectional connections between

end user equipments. An egress LSR may be adapted to select a reservation style for a RSVP session from a set of possible reservation styles, a RSVP session being identified by a unique tuple, and a RSVP session may be adapted to create one, or more, LSPs, depending on the selected reservation style. The set of reservation styles may include, inter alia, Fixed Filter (FF), Wildcard Filter (WF) and Shared Explicit (SE).

The LSPs may be set up by egress LSRs, and an Open Short Path First (OSPF) routing infrastructure of said system may be adapted to indicate said LSRs.

The egress LSRs may be provided by either Area Border Routers (ABRs), or Autonomous System Border Routers (ASBRs), or edge LSRs with attached subnets, and the LSRs may be adapted to advertise the different types of routes that define Forward Equivalence Class (FEC) elements. ABR nodes may be adapted to generate FECs that contain address prefixes for at least inter-area routes, or intra-area routes, and edge LSR nodes may be adapted to generate FECs for their attached subnets. The edge LSR nodes may be adapted to define one FEC for all of the address prefixes of their attached subnets.

The egress LSRs may be adapted to use Label Mapping (LM) to distribute FEC labels to upstream neighbours.

According to a second aspect of the present invention, there is provided, a method for establishing end-to-end data transmission paths in a data transmission system including a backbone network, a plurality of routers, and a plurality of end user equipments, said method being characterised by said end user equipments using RSVP messages to convey their specific requirements and by said system using MPLS to intercept said RSVP messages and establish Label Switched Paths (LSPs) over said backbone network, based on the content of said RSVP messages.

The method may be characterised by said system providing end-to-end Quality of Service (QoS) through interworking between RSVP and MPLS.

The method may be characterised by said system aggregating a number of RSVP messages, having the same QoS semantics, into a single LSP.

The method may be characterised by said LSPs being established between ingress and egress LSRs of said backbone network, and by said end user equipments being connected to said LSRs via end user networks to which said end user equipments are connected.

The method may be characterised by an end user equipment (TE1) requiring a connection to another end user equipment (TE2) issuing an RSVP-PATH message, indicating a required QoS, to an ingress LSR of said backbone network; via a first end user network with which said end user equipment (TE1) is associated; said ingress LSR encapsulating said RSVP-PATH message and sending it through said backbone network to an egress LSR of said backbone network; said egress LSR, on receipt of said RSVP-PATH message, decapsulating the message and issuing a PATH message towards a second end user network with which said end user equipment (TE2) is associated; said end user equipment (TE2) on acceptance of the connection to end user equipment (TE1) sending an RESV message towards said end user equipment (TE 1); said egress LSR, upon receipt of the RESV message: buffering the RESV message, and establishing a Label Switched Path (LSP) to said ingress router, based on the parameters contained in the RESV message; said egress LSR, upon establishment of the LSP, encapsulating the buffered RESV message and sending it via the established LSP to said ingress LSR; said ingress LSR, on receipt of said RESV message, decapsulating the RESV message and sending it towards the end user equipment (TE1), thereby establishing a unidirectional connection from end user equipment (TE1) to end user equipment (TE2).

The method may be characterised by an egress LSR of said backbone network selecting a reservation style for a RSVP session from a set of possible reservation styles, a RSVP session being identified by a unique tuple, and by said RSVP session creating one, or more, LSPs, depending on the selected reservation style.

The method may be characterised by said set of reservation styles including, inter alia, Fixed Filter (FF), Wildcard Filter (WF) and Shared Explicit (SE).

The method may be characterised by egress LSRs, setting up LSPs, and by an Open Short Path First (OSPF) routing infrastructure of said system indicating these nodes.

The method may be characterised by said egress LSRs being provided by either Area Border Routers (ABRs), or Autonomous System Border Routers (ASBRs), or edge LSRs with attached subnets, and by said LSRs advertising the different types of routes that define Forward Equivalence Class (FEC) elements.

The method may be characterised by said ABR nodes generating FECs that contain address prefixes for at least inter-area routes, or intra-area routes, and by said edge LSR nodes generating FECs for their attached subnets.

The method may be characterised by said edge LSR nodes defining one FEC for all of the address prefixes of their attached subnets.

The method may be characterised by said egress LSRs using Label Mapping (LM) to distribute FEC labels to upstream neighbours.

The method may be characterised by each LSR on a path towards an ingress LSR, on receipt of a LM message checking whether, or not, an advertiser according to routing is the next hop for a FEC; and either withdrawing a downstream label, in the event that it is not the next hop, by responding with a Label Withdraw (LW) message, or checking whether its router-id is in the LSR-path-vector, in the event that it is the next hop.

The method may be further characterised by each of said LSRs discarding the label mapping, in the event that said router-id is in said vector, and sending a NAK to the advertiser; or, in the event that said router-id is not in said vector, finding an upstream label for the FEC, cross-connecting it with the downstream label, and sending a Label Mapping (LM) message for the upstream label to all upstream

neighbours, said LM message having the LSR's router-id added to the LSR-path-vector list, and the hop count set to a value of the received LM message plus 1.

The method may be characterised by said ingress LSRs, on receipt of LMs retaining only the labels used for forwarding; and investigating an LSR-path-vector list to check whether a loop has been established.

The method may be characterised by said ingress LSRs, on receipt of an IP data packet, modifying the TTL value based on the hop count value associated with a FEC; attaching a shim header to the IP data packet; and sending it on a Virtual Channel (VC) associated to the FEC. The method may be further characterised by said shim header having a label of value 0, indicating that forwarding of said IP data packet be based on a IPv4 header, and carrying a TTL value equal to a modified TTL of said IP header; and by a CoS (Class of Service) value being set to best effort. The method may be further characterised by said egress LSRs, on receipt of a labelled IP data packet, removing said shim header and delivering the IP data packet to an IP layer for further forwarding.

The method may be characterised by periodically refreshing said RSVP, PATH and RESV message

The foregoing and other features of the present invention will be better understood from the following description with reference to the accompanying drawings, in which:

Figure 1 diagrammatically illustrates the interworking between RSVP and MPLS, according to the present invention; and

Figure 2 diagrammatically illustrates egress based LSP initialisation used by a data transmission system of the present invention.

A glossary of the abbreviations used in this patent specification is set out below to facilitate an understanding of the present invention.

ABR:	Area Border Router
ASBR:	Autonomous System Border Router
ATM:	Asynchronous Transfer Mode
BGP4:	Border Gateway Protocol Version 4
CoS:	Class of Service
FEC:	Forwarding Equivalence Class
FIB:	Forwarding Information Base
FF:	Fixed Filter
IETF:	Internet Engineering Task Force
IP:	Internet Protocol
IPv4:	Internet Protocol Version 4
LDP:	Label Distribution Protocol
LIB:	Label Information Base
LM:	Label Mapping
LR:	Label Request
LSP:	Label Switched Path
LSR:	Label Switched Router
LW:	Label Withdraw

NAK: Negative Acknowledgement

OSPF: Open Short Path First

QoS: Quality of Service

RSVP: Resource Reservation Protocol

SE: Shared Explicit

TE: Terminal Equipment

TTL: Time To Live

TLV: Type Length Value

UBR: Unspecified Bit Rate (ATM)

VC: Virtual Circuit

WF: Wildcard Filter

It will be seen from the subsequent description that the present invention provides a system architecture and protocol enabling end users to convey their specific end-to-end QoS requirements by means of connecting RSVP islands through a MPLS cloud. End-users convey their requirements through use of RSVP messages. MPLS intercepts the RSVP messages and establishes Label Switched Paths (LSPs) over a backbone network (e.g. ATM network) based on the information indicated in the RSVP messages. MPLS also aggregates RSVP flows (depending on the reservation style) so that the number of flow states are substantially reduced. In the case of conveyed QoS requirements, this gives rise to an end-to-end QoS:d network that scales much better than an RSVP only network. Using MPLS alleviates the RSVP scalability problems, referred to above, since individual RSVP flows are

aggregated in a backbone network, such as, an ATM network, and switched at layer 2 rather than being processed at layer 3 hop-by-hop. The major advantage of layer-2 switching is that it boosts performance as opposed to conventional routers. Furthermore, MPLS LSPs are established based on the true end-user requirements rather than an artificial classification procedure in the edge routers as it is proposed today.

Figure 1 of the accompanying drawings diagrammatically illustrates the manner in which interworking is effected between RSVP and MPLS. As illustrated in Figure 1, end user Terminal Equipments, TE1 and TE2, are respectively connected to separate networks, for example, Local Area Networks (LANs), which form the RSVP islands, referred to above. In practice, a number of TEs would be connected to each of the LANs. The MPLS cloud may, for example, be provided by a backbone network including a number of interconnected switches (e.g. ATM switches) connected to Label Switched Routers, such as, LSR1 and LSR2, at the edge of the backbone network. The Terminal Equipments, TE1 and TE2 are, therefore, respectively connected to LSR1 and LSR2 via a LAN and thereby to the backbone network.

In MPLS, the establishment of a label switched path (LSP) is primarily based on the destination network prefix. However, there is also a possibility of classifying data packets based on the source address, or the destination address, or the port-id, or the transport protocol, etc., or any combination thereof (a.k.a flow). It is considered that a destination-network-prefix based classification is useful for aggregating best-effort flows while flow specific LSPs are useful for time critical applications, or applications for which specific guarantees are requested. However, classifying data packets based on transport layer information may not be the best way of estimating user requirements. It is reasonable to assume that end users will use RSVP to convey their specific requirements. It is also reasonable to assume that MPLS will play an important role in the backbone. Thus, connecting RSVP islands through a MPLS cloud, as is illustrated in Figure 1 of the accompanying drawings, becomes a key issue for enabling end-to-end guarantees.

The overall operation of the Resource Reservation Protocol (RSVP) will now

be described, by way of example, with reference to Figure 1 of the accompanying drawings. Thus, operation of RSVP involves the following procedural steps:

1. The Terminal Equipment, TE l, issues an RSVP-PATH message indicating the QoS requirements via the LAN (RSVP island) to the ingress edge router, LSR1.

2. The ingress edge router, LSR1, encapsulates the RSVP-PATH message and sends it through the MPLS cloud, i.e. through the backbone network, to the egress edge router, LSR2.

3. The egress edge router, LSR2, decapsulates the RSVP-PATH message and issues the PATH message towards the LAN forming part of the receiving the RSVP island.

4. The Terminal Equipment, TE2, if it accepts the connection, sends an RESV message towards the Terminal Equipment, TE 1.

5. Upon receipt of the RESV message, the edge router, LSR2, buffers the RESV message and establishes a Label Switched Path (LSP) to the edge router, LSR1, based on the parameters contained in the RESV message; the procedures used to establish the LSPs will be outlined in the subsequent description.

6. On the establishment of a LSP, the edge router, LSR2, encapsulates the previously buffered RESV message and sends it via the established LSP to the edge router, LSR1

7. On receipt of the RESV message, the edge router, LSR1, decapsulates the RESV message and sends it towards the Terminal Equipment, TE1, via the associated LAN.

At the end of the foregoing procedure, a unidirectional connection will be established from the Terminal Equipment, TE1 to the Terminal Equipment, TE2.

If the Terminal Equipment, TE2, would like to set up a connection to the Terminal Equipment, TE1, the same procedure, as outlined above, would be executed but in the reverse direction, i.e. TE2 would issue the RSVP-PATH message, the edge router, LSR2 would be the ingress router and the edge router, LSR1, would be the egress router.

The manner in which aggregation of RSVP flows to Label Switched Paths (LSPs) is effected will now be described. Since LSRs constitute an aggregation point, several RSVP flows can be aggregated into the same LSP as long as they have the same QoS semantics. However, a receiver node can select from among a set of possible reservation styles for each session, and each RSVP session must have a particular style. In other words, senders have no influence on the choice of reservation style. The receiver can choose different reservation styles that get mapped to different LSPs.

An RSVP session is identified by a unique (destination address, protocol, destination port) tuple. Thus, an RSVP session can create one, or more, LSPs, depending on the reservation style which is chosen. Some reservation styles, such as Fixed Filter (FF), dedicate a particular reservation to an individual sender node. Other reservation styles, such as Wildcard Filter (WF) and Shared Explicit (SE), can share a reservation among several sender nodes.

The Fixed Filter (FF) style of reservation creates a distinct reservation for traffic from each sender that is not shared by other senders. This style is common for applications in which traffic from each sender is likely to be concurrent and independent. The total amount of reserved bandwidth on a link for sessions using FF style is the sum of the reservations for the individual senders. Because each sender has its own reservation, a unique label and a separate label switched path is assigned to each sender. This results in a point-to-point LSP between every sender-receiver pair.

With the Wildcard Filter (WF) style of reservation, a single shared reservation is used for all senders. The total reservation on a link remains the same regardless of the number of senders. This style is useful in applications in which not all senders

send traffic at the same time. A single label-switched-path is created for all senders, because all senders to the session are covered by the reservation. On links that senders share, a single label is allocated. If there is only one sender, the LSP looks like a normal point-to-point connection. When multiple senders are present, a multipoint-to-point LSP is created. This has the advantage of minimizing the number of LSPs allowing better scaling of the network.

With the Shared Explicit (SE) style of reservation, any sender is allowed to share the reservation, the SE style allowing a receiver to explicitly specify the senders to be included. There is a single reservation on a link for all of the senders listed in the RSVP message. Only listed senders can join the reservation. Because each sender is explicitly listed in the RESV message, separate labels can be assigned to each sender, thereby creating separate LSPs for each sender. Unlike FF, all SE LSPs share a single reservation. Unlike WF, a separate LSP is created for each sender.

The procedures used to establish egress initiated Label Switched Paths (LSPs) will now be described. These procedures are based on a combination of an independent label distribution and a conservative label retention mode. In particular, Label Switched Paths (LSPs) are set-up by egress Label Switched Routers (LSRs), for example, Area Border Routers (ABRs), or Autonomous System Border Routers (ASBRs), or edge LSRs with attached subnets. The Open Short Path First (OSPF) routing infrastructure indicates these nodes. The different types of routes advertised by these nodes define the Forward Equivalents Class (FEC) elements of the FEC TLV and thus achieving stream aggregation. For example, an ABR defines a FEC containing address prefixes for summary (inter-area), intra-area, or other types of routes. ASBRs build a FEC for the external routes redistributed by BGP4. Edge LSRs generate FECs for their attached subnets (intra-area routes). One FEC can be defined for all the address prefixes of attached subnets.

Each egress LSR distributes the labels for its FECs to all its upstream neighbours through LDP. Message Label Mapping (LM) is used for that purpose. Egress based LSP initialisation is diagrammatically illustrated, by way of example, in Figure 2 of the accompanying drawings.

As illustrated in Figure 2, Area Border Router, ABR1, sends to its neighbour, Label Switched Router, LSR1, a Label Mapping (LM) for the address prefixes 'n' and 'r' defined in one FEC, i.e. Label: X, FEC (n,r). For the sake of loop detection, the router-id of the LSR is put in the LSR-path-vector TLV of the message. Finally, the Hop Count TLV of the message is set to 1. The Hop Count is used by the ingress LSR, LSR2 of Figure 2, to calculate the TTL (Time to Live) value of the IP data packet entering the MPLS backbone network (e.g. ATM network).

In practice, each core LSR on the path towards ingress LSRs will, upon receiving a Label Mapping message, follow the procedural steps set out below.

Step 1: Check if the advertiser according to routing is the next hop for the FEC. If it is not, the downstream label is withdrawn through responding with a Label Withdraw (LW) message - see the response of LSRI to ABR3 in Figure 2. If it is the next hop, proceed to Step 2.

Step 2. Check if its router-id is in the LSR-path-vector, i.e. a loop has been established. If yes, discards the label mapping and sends a NAK to the advertiser. Otherwise, proceed to Step 3.

Step 3. Find an upstream label for the FEC, cross-connect it with the downstream label, and send a Label Mapping (LM) message for the upstream label to all upstream neighbours. The LM message will have the LSR's router-id added to the LSR-path-vector list, and the Hop Count set to the value of the received LM message plus 1.

The establishment of upstream and downstream labelling is simplified if VC (Virtual Circuit) merging is supported by the LSRs. Otherwise, each LSR has to request a downstream label for each upstream label. This results in an overhead in LDP signalling and state information.

Ingress LSRs, upon receiving Label Mappings (LMs), retain only the labels used for forwarding. Ingress LSRs check whether a loop has been established by investigating the LSR-path-vector list. These labels and their FECs, as well as the

Hop Count values, will populate the LIB and become the default paths for best effort traffic.

Upon receipt of an IP data packet, an ingress Label Switched Router, i.e. LSR2 of Figure 2:

- modifies the TTL value based on the Hop Count value associated with the FEC;
- attaches a shim header to the IP data packet; and
- sends it on the VC associated to the FEC.

The shim header carries the label of value 0 (IPv4 Explicit Null Label) indicating that the label stack must be popped at the egress LSR and the forwarding of the IP data packet should be based on the IPv4 header. It also carries a TTL value equal to the modified TTL of the IP header. The CoS (Class of Service) value is set to best effort.

When a labelled IP data packet is received by an egress LSR, the shim header is removed and the packet is delivered to the IP layer for further forwarding.

Upon routing changes, downstream labels for the affected FECs are released if they are not in use for other FEC elements. New labels are requested from the new next hop. Finally, new label mappings upstream may be required.

The tables shown in Figure 2 provide further information concerning the FECs, Labels, Hop Count, Address Prefixes, and Next Hop for Label Switched Routers, LSR1 and LSR2, together with VC lookup tables for LSR1.

The main advantage of the egress based solution is its economic use of VCs due to the mapping of a number of FECs onto one LSP, i.e. stream aggregation. It should, however, be noted that stream aggregation increases the traffic load on VCs, which may result in low performance if per VC queuing, due to fairness, is used for

UBR. Another downside is protocol complexity. Loops might arise but the system is adapted to detect and break them by means of the LSR-vector-path.

In RSVP, PATH and RESV message are periodically refreshed. Supporting a large number of sessions may present scaling problems. Since memory and processing requirements increase linearly with an increase in the number of states, one way to alleviate the scaling problem is to increase the Refresh timer. However, this is done at the cost of increasing refresh timeout. It is, however, possible to increase the refresh timeout and still be able to react for faster detection of connectivity problems. For example, as soon as topology failures occur, every node adjacent to the failure notifies all affected nodes.

CLAIMS

1. A data transmission system including a backbone network, a plurality of routers, and a plurality of end user equipments, said system being adapted to establish end-to-end data transmission paths over the backbone network, characterised in that said end user equipments are adapted to use RSVP messages to convey their specific requirements and in that said system is adapted to use MPLS to intercept said RSVP messages and establish Label Switched Paths (LSPs) over said backbone network, based on the content of said RSVP messages.

2. A data transmission system, as claimed in claim 1, characterised in that said system is adapted to provide end-to-end Quality of Service (QoS) through interworking between RSVP and MPLS.

3. A data transmission system, as claimed in claim 2, characterised in that said system is adapted to aggregate a number of RSVP messages, having the same QoS semantics, into a single LSP.

4. A data transmission system, as claimed in any preceding claim, characterised in that said system further includes a plurality of end user networks, to which said end user equipments are adapted to be connected, and a plurality of routers adapted to connect said end user networks to said backbone network.

5. A data transmission system, as claimed in claim 4, characterised in that at least some of said end user networks are Local Area Networks (LANs).

6. A data transmission system, as claimed in claim 4, or claim 5, characterised in that said backbone network is an ATM network, adapted for the transmission of IP data packets and including a number of interconnected ATM switches, in that at least two of said routers are edge routers for said ATM network, and in that said edge routers are Label Switched Routers (LSRs), adapted to be connected to end user equipments for the establishment of unidirectional connections between end user equipments.

7. A data transmission system, as claimed in claim 6, characterised in that an egress LSR is adapted to select a reservation style for a RSVP session from a set of possible reservation styles, a RSVP session being identified by a unique tuple, and in that a RSVP session is adapted to create one, or more, LSPs, depending on the selected reservation style.

8. A data transmission system, as claimed in claim 7, characterised in that said set of reservation styles includes, inter alia, Fixed Filter (FF), Wildcard Filter (WF) and Shared Explicit (SE).

9. A data transmission system, as claimed in either claim 7, or claim 8, characterised in that LSPs are set up by egress LSRs, and in that an Open Short Path First (OSPF) routing infrastructure of said system is adapted to indicate said LSRs.

10. A data transmission system, as claimed in claim 9, characterised in that said egress LSRs are provided by either Area Border Routers (ABRs), or Autonomous System Border Routers (ASBRs), or edge LSRs with attached subnets, and in that said LSRs are adapted to advertise the different types of routes that define Forward Equivalence Class (FEC) elements.

11. A data transmission system, as claimed in claim 10, characterised in that ABR nodes are adapted to generate FECs that contain address prefixes for at least inter-area routes, or intra-area routes, and in that edge LSR nodes are adapted to generate FECs for their attached subnets.

12. A data transmission system, as claimed in claim 11, characterised in that said edge LSR nodes are adapted to define one FEC for all of the address prefixes of their attached subnets.

13. A data transmission system, as claimed in any of claims 7 to 12, characterised in that said egress LSRs are adapted to use Label Mapping (LM) to distribute FEC labels to upstream neighbours.

14. A method for establishing end-to-end data transmission paths in a data

transmission system including a backbone network, a plurality of routers, and a plurality of end user equipments, said method being characterised by said end user equipments using RSVP messages to convey their specific requirements and by said system using MPLS to intercept said RSVP messages and establish Label Switched Paths (LSPs) over said backbone network, based on the content of said RSVP messages.

15. A method, as claimed in claim 14, characterised by said system providing end-to-end Quality of Service (QoS) through interworking between RSVP and MPLS.

16. A method, as claimed in claim 15, characterised by said system aggregating a number of RSVP messages, having the same QoS semantics, into a single LSP.

17. A method, as claimed in any of claims 14 to 16, characterised by said LSPs being established between ingress and egress LSRs of said backbone network, and by said end user equipments being connected to said LSRs via end user networks to which said end user equipments are connected.

18. A method, as claimed in claim 17, characterised by:

- an end user equipment (TE1) requiring a connection to another end user equipment (TE2) issuing an RSVP-PATH message, indicating a required QoS, to an ingress LSR of said backbone network; via a first end user network with which said end user equipment (TE1) is associated;

- said ingress LSR encapsulating said RSVP-PATH message and sending it through said backbone network to an egress LSR of said backbone network;

- said egress LSR, on receipt of said RSVP-PATH message, decapsulating the message and issuing a PATH message towards a second end user network with which said end user equipment (TE2) is associated;

- said end user equipment (TE2) on acceptance of the connection to end user

equipment (TE1) sending an RESV message towards said end user equipment (TE 1);

- said egress LSR, upon receipt of the RESV message:

  - buffering the RESV message; and

  - establishing a Label Switched Path (LSP) to said ingress router, based on the parameters contained in the RESV message;

- said egress LSR, upon establishment of the LSP, encapsulating the buffered RESV message and sending it via the established LSP to said ingress LSR;

- said ingress LSR, on receipt of said RESV message, decapsulating the RESV message and sending it towards the end user equipment (TE1), thereby establishing a unidirectional connection from end user equipment (TE1) to end user equipment (TE2).

19. A method, as claimed in either claim 17, or claim 18, characterised by an egress LSR of said backbone network selecting a reservation style for a RSVP session from a set of possible reservation styles, a RSVP session being identified by a unique tuple, and by said RSVP session creating one, or more, LSPs, depending on the selected reservation style.

20. A method, as claimed in claim 19, characterised by said set of reservation styles including, inter alia, Fixed Filter (FF), Wildcard Filter (WF) and Shared Explicit (SE).

21. A method, as claimed in any of claims 17 to 20, characterised by egress LSRs, setting up LSPs, and by an Open Short Path First (OSPF) routing infrastructure of said system indicating these nodes.

22. A method, as claimed in claim 21, characterised by said egress LSRs being provided by either Area Border Routers (ABRs), or Autonomous System Border Routers (ASBRs), or edge LSRs with attached subnets, and by said LSRs

advertising the different types of routes that define Forward Equivalence Class (FEC) elements.

23. A method, as claimed in claim 22, characterised by said ABR nodes generating FECs that contain address prefixes for at least inter-area routes, or intra-area routes, and by said edge LSR nodes generating FECs for their attached subnets.

24. A method, as claimed in claim 23, characterised by said edge LSR nodes defining one FEC for all of the address prefixes of their attached subnets.

25. A method, as claimed in any of claims 17 to 24, characterised by said egress LSRs using Label Mapping (LM) to distribute FEC labels to upstream neighbours.

26. A method, as claimed in claim 25, characterised by each LSR on a path towards an ingress LSR, on receipt of a LM message:

- checking whether, or not, an advertiser according to routing is the next hop for a FEC; and either:

  - withdrawing a downstream label, in the event that it is not the next hop, by responding with a Label Withdraw (LW) message; or

  - checking whether its router-id is in the LSR-path-vector, in the event that it is the next hop.

27. A method, as claimed in claim 26, characterised by each of said LSRs:

- discarding the label mapping, in the event that said router-id is in said vector, and sending a NAK to the advertiser; or

- in the event that said router-id is not in said vector:

  - finding an upstream label for the FEC;

  - cross-connecting it with the downstream label; and

- sending a Label Mapping (LM) message for the upstream label to all upstream neighbours, said LM message having the LSR's router-id added to the LSR-path-vector list, and the hop count set to a value of the received LM message plus 1.

28. A method, as claimed in either claim 26 or claim 27, characterised by said ingress LSRs, on receipt of LMs:

- retaining only the labels used for forwarding; and

- investigating an LSR-path-vector list to check whether a loop has been established.

29. A method, as claimed in any of claims 26 to 28, characterised by said ingress LSRs, on receipt of an IP data packet:

- modifying the TTL value based on the hop count value associated with a FEC;

- attaching a shim header to the IP data packet; and

- sending it on a Virtual Channel (VC) associated to the FEC.

30. A method, as claimed in claim 29, characterised by said shim header:

- having a label of value 0, indicating that forwarding of said IP data packet be based on a IPv4 header;

- carrying a TTL value equal to a modified TTL of said IP header;

and by a CoS (Class of Service) value being set to best effort.

31. A method, as claimed in either claim 29, or claim 30, characterised by said egress LSRs, on receipt of a labelled IP data packet, removing said shim header and

delivering the IP data packet to an IP layer for further forwarding.

32. A method, as claimed in any of claims 17 to 31, characterised by periodically refreshing said RSVP, PATH and RESV message

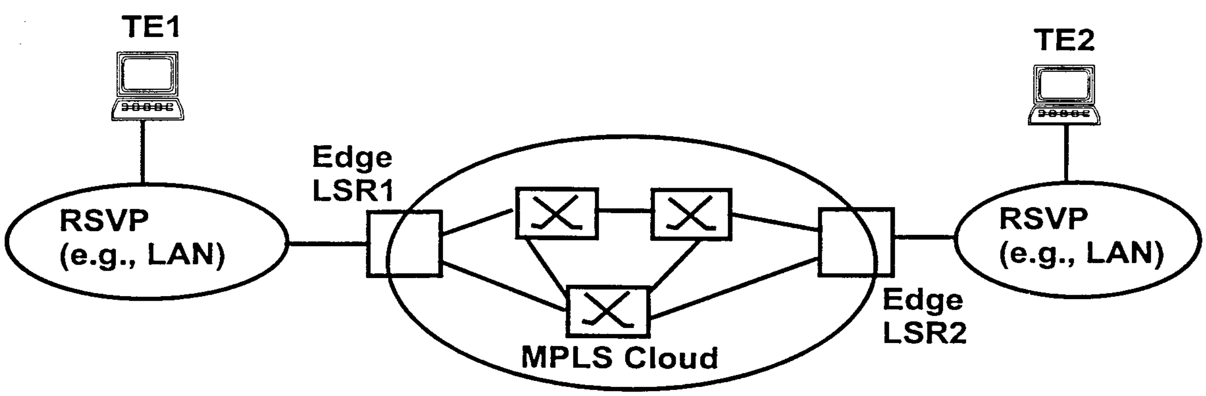

FIGURE 1

LSR2:s LIB

FEC	Label	Hop Count
n	Z	2
r	E	2
m	E	2

LSR1:s FIB

Address Prefix	Next Hop
n	ABR1
r	ABR2
m	ABR3

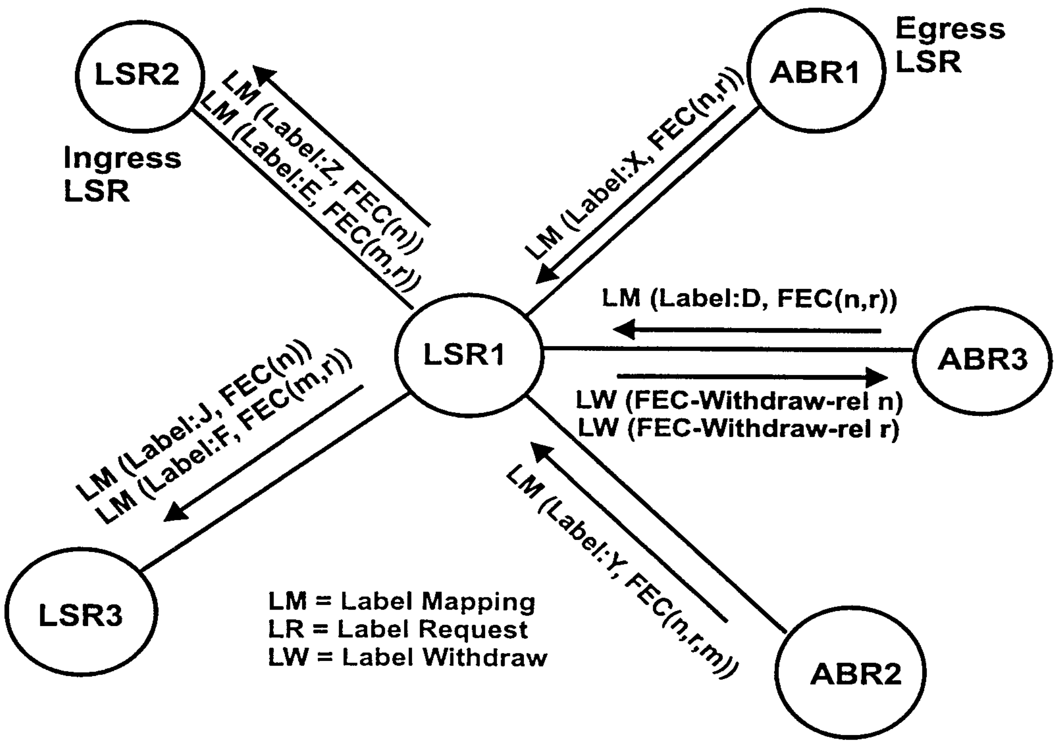

LSR1:s LIB

FEC	Label
n	X
r	Y
m	Y

LSR1:s VC Lookup Table

VCin	VCout
Z	X
F	Y
E	Y
J	X

FIGURE 2

INTERNATIONAL SEARCH REPORT

International application No.
PCT/SE 99/02278

A. CLASSIFICATION OF SUBJECT MATTER

IPC7: H04Q 11/04, H04L 12/56
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC7: H04L, H04Q, H04J

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	VISWANATHAN, et al. "Soft State Switching A Proposal to Extend RSVP for Switching RSVP Flows" (online), August 1997. (retrieved on 2000-04-15). Retrieved from the Internet: <URL: http://infonet.aist-nara.ac.jp/member/nori-d/mlr/id/draft-viswanathan-mpls-rsvp-00.txt>	1-4,6,14-17
Y		5,7-13,18-32
	--	
Y	DAVIE, et al. "Use of Label Switching With RSVP" (online), March 1998. (retrieved on 2000-04-15). Retrieved from the Internet: <URL: http://infonet.aist-nara.ac.jp/member/nori-d/mlr/id/draft-ietf-mpls-rsvp-00.txt>	5,18
A		1-4,6-17, 19-32
	--	

[X] Further documents are listed in the continuation of Box C. [X] See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" erlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
17 April 2000	26-04-2000
Name and mailing address of the ISA/ Swedish Patent Office Box 5055, S-102 42 STOCKHOLM Facsimile No. +46 8 666 02 86	Authorized officer Erik Johannesson/CL Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.
PCT/SE 99/02278

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	SWALLOW, et al. "Extensions to RSVP for LSP Tunnels" (online), November 1998. (retrieved on 2000-04-15). Retrieved from the Internet: <URL: http:/infonet.aist-nara.ac.jp/member/nori-d/mlr/id/draft-ietf-mpls-rsvp-tunnel-00.txt>	7-13,19-32
A		1-6,14-18
	--	
Y	DAVIE, et al. "Use of Label Switching With ATM" (online), July 1998. (retrieved on 2000-04-15). Retrieved from the internet: <URL: http://ardnoc41.canet2.net/mpls/drafts/draft-davie-mpls-atm-01.txt>	10-13,22-32
A		1-9,14-21
	--	
A	GB 2320159 A (INTERNATIONAL BUSINESS MACHINES CORPORATION), 10 June 1998 (10.06.98), see whole document	1-32
	--	
A	EP 0790751 A2 (LUCENT TECHNOLOGIES INC.), 20 August 1997 (20.08.97), see whole document	1-32
	--	

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/SE 99/02278

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
GB 2320159 A	10/06/98	GB 9720770 D	00/00/00
EP 0790751 A2	20/08/97	CA 2196006 A	17/08/97
		JP 9247190 A	19/09/97